

Vehicle routing by learning from historical solutions (extended abstract)*

Rocsildes Canoy and Tias Guns

Vrije Universiteit Brussel, Brussels, Belgium
{Rocsildes.Canoy,Tias.Guns}@vub.be

Abstract

An important initial step in solving vehicle routing problems (VRP) is the careful formulation of the problem objectives and constraints. Oftentimes in practice, the optimization of the route plans takes into account not only time and distance-related factors, but also a variety of constraints related to inventory and scheduling, environmental and energy concerns, personal preferences of route planners and drivers, etc. As each of these factors are introduced in the formulation, the VRP becomes more and more complex. Furthermore, besides the added complexity, there is no guarantee that the solution resulting from the formulation would satisfy the route planners and all involved stakeholders. Consequently, route planners end up having to spend additional time and effort in modifying the VRP solution until the desired result is obtained.

Proposed Method. In the paper, we present a novel approach to solving the VRP which does not require explicit problem characterization. Instead, by assuming the existence of past solutions over similar sets of customers, we propose a learning method that can generate results which closely approximate the previous solutions and would hence require fewer manual modifications. Specifically, by using a first-order Markov model, our method learns a probability transition matrix from historical solutions to predict the routes for an entire fleet.

To learn from historical data, we take inspiration from various machine learning papers on route prediction *for a single vehicle*. Markov models developed from historical data have been applied to driver turn prediction, prediction of the remainder of the route by looking at previous road segments taken by the driver, and predicting individual road choices given the origin and destination. These studies have produced positive and encouraging results. Hence, in this work, we investigate the use of Markov models for predicting the route choices *for an entire fleet*, and how to use these choices to solve the VRP.

In the paper, we describe in detail our algorithm for constructing the probability transition matrix from historical data. Each instance of the historical data is a route plan over a set of customers. Instances are ordered over time, and

* Copyright ©2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).




Fig. 1. Route Difference




Fig. 2. Arc Difference

the set of customers can vary from instance to instance. To account for this, we additionally propose three weighing schemes to combine the Markov models that we built for each historic instance separately. The three schemes are: UNIF, where weights are uniformly distributed; TIME for time-based weighing, where older instances are given smaller weights; and SIMI, where weights are adjusted according to the similarity of the customer sets. Squared weights, TIME2 and SIMI2, are added in the numerical experiments.

Numerical Experiments. The algorithm was implemented in Python 3.6.5 with the CPLEX 12.8 solver. We used real-life data of 201 instances (days), grouped by day-of-week provided by a small transportation company. They have an average of 8.7 vehicles servicing 35.1 stops per instance. We conducted experiments to compare the performances of our proposed weighing schemes against the classical distance-based (DIST) solution method. Route prediction accuracy was measured by the number of customers incorrectly assigned to the wrong route (Route Difference) and the number of arcs taken in the predicted solution but not in the actual plan (Arc Difference).

Results show that all the proposed schemes consistently outperformed DIST (Figs. 1 and 2). A slight improvement in prediction accuracy can also be observed when using squared weights, e.g., TIME vs. TIME2. Among all the proposed schemes, TIME2 gave the most accurate predictions.

Along with a visual example showing how the routes are predicted by the transition matrix, several other experimental results are presented in the paper. Altogether, these results have confirmed the ability of the proposed method to learn the solution structure. An added advantage is that, due to the characteristic sparsity of the transition matrix, computation is found to be much faster and more efficient than the traditional distance-based method (a minute versus > 10 minutes to compute the optimal solution).

The full paper is published as:

Rocsildes Canoy and Tias Guns: Vehicle routing by learning from historical solutions. Principles and Practice of Constraint Programming, 25th International Conference (CP 2019), pages 1–16, Springer (2019).